EXHIBIT 9

An Owner's Manual for the Internet

Includes CD-ROM

Internet Core Protocols

The Definitive Guide

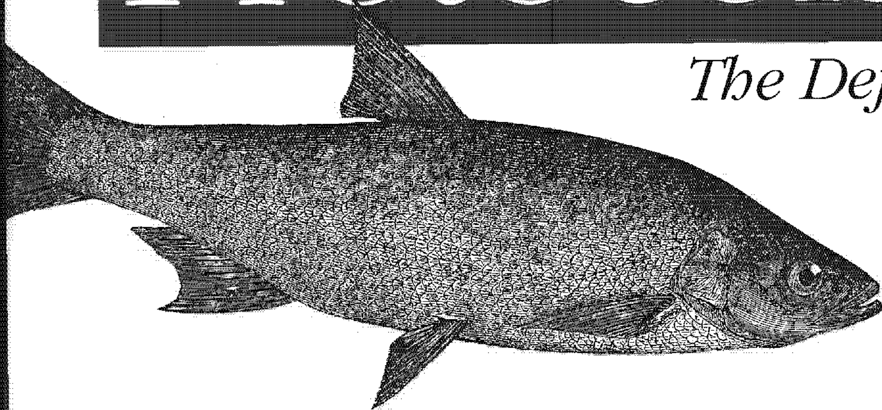

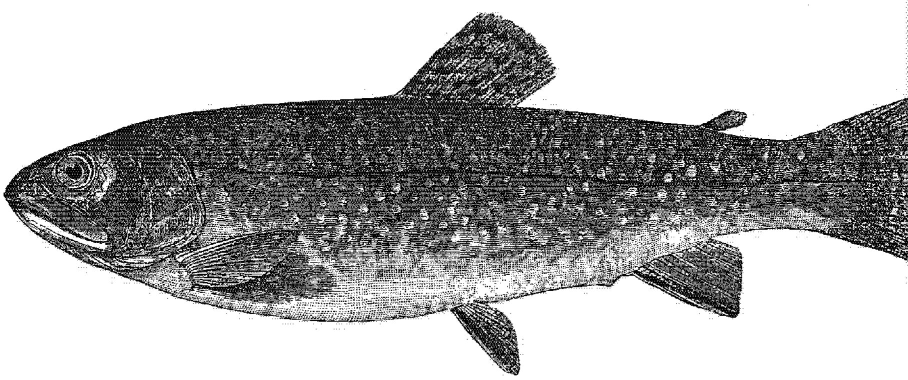

Eric A. Hall
Foreword by Vint Cerf

O'REILLY®

Internet Core Protocols: The Definitive Guide
by Eric A. Hall

Published by O'Reilly & Associates, Inc., 101 Morris Street, Sebastopol, CA 95472.

Editor: Mike Loukides

Production Editor: Nicole Arigo

Cover Designer: Edie Freedman

Printing History:

February 2000: First Edition.

Library of Congress Cataloging-in-Publication Data:

Hall, Eric A.
Internet core protocols : the definitive guide / Eric A. Hall.
p. cm.
Includes bibliographical references.
ISBN 1-56592-572-6 (alk. paper)
1. Internet (Computer network) 2. Computer network protocols. I. Title.

TK5105.875.I57 H355 2000
004.6'2--dc21

99-089559

ISBN: 1-56592-572-6
[M] [7/02]

1

An Introduction to TCP/IP

If you've been using TCP/IP-based networking products for any length of time at all, you're probably already aware of how IP addressing, routing, and other fundamental aspects of the Internet family of protocols work, at least from a user's perspective.

What you probably don't know—unless you've been formally trained in these subjects—is what makes TCP/IP work from the wire's perspective, or from the perspective of the applications in use on your network. This chapter provides you with an introduction to these viewpoints, providing you with a better understanding of the nature of the traffic on your network.

A Brief History of the Internet

Before you can understand how TCP/IP works—or why it works the way it does—you first have to understand the origins of the networking protocols and the history of the Internet. These subjects provide a foundation for understanding the basic design principles behind TCP/IP, which in turn dictate how it is used today.

TCP/IP presented a radical departure from the traditional computer networking services in use during its development. In the early days of commercial computing (the late 1960s), most companies bought a single large system for all of their data processing needs. These systems used proprietary networking architectures and protocols, which primarily consisted of plugging dumb terminals or line printers into an intelligent communications controller, each of which used proprietary networking protocols to communicate with the central hosts.

Most of the early computer networks used this hierarchical design for their proprietary network protocols and services. As users' computing requirements expanded,

they rarely bought a different system from a different vendor, but instead added new components to their existing platforms or replaced the existing system with a newer, larger model. Cross-platform connectivity was essentially unheard of, and was not expected. To this day, you still can't plug an IBM terminal into a DEC system and expect it to work. The protocols in use by those devices are completely different from each other.

As the use of computers became more critical to national defense, it became clear to the U.S. military in particular that major research centers and institutions needed to be able to share their computing resources cooperatively, allowing research projects and supercomputers to be shared across organizational boundaries. Yet, since each site had different systems (and therefore different networking technologies) that were incompatible with the others, it was not possible for users at one site to use another organization's computing services easily. Nor could programs easily be ported to run on these different systems, as each of them had different languages, hardware, and network devices.

In an effort to increase the sharing of resources, the Advanced Research Projects Agency (ARPA) of the Department of Defense (DOD) began coordinating the development of a vendor-independent network to tie the major research sites together. The need for a vendor-independent network was the first priority, since each facility used different computers with proprietary networking technology. In 1968, work began on a private packet-switched network, which eventually became known as ARPAnet.

ARPAnet was the world's first wide-area packet-switching network, designed to allow individual units of data to be routed across the country as independent entities. Previous networks had been circuit-switched, involving dedicated end-to-end connections between two specific sites. In contrast, the ARPAnet allowed organizations to interconnect into a mesh-like topology, allowing data to be sent from one site to another using a variety of different routes. This design was chosen for its resilience and built-in fault-tolerance: if any one organization were bombed or otherwise removed from the network, it wouldn't affect the rest of the organizations on the network.

During this same time period, other network providers also began interconnecting with the ARPAnet sites, and when these various networks began connecting to each other, the term "Internet" came into use. Over the next few years, more organizations were added to the ARPAnet, while other networks were also being developed, and new network technologies such as Ethernet were beginning to gain popularity as well.

All of this led to the conclusion that networking should be handled at a higher layer than was allowed by the ARPAnet's packet-switching topology. It became

increasingly important to allow for the exchange of data across different physical networks, and this meant moving to a set of networking protocols that could be implemented in software on top of any physical topology, whether that be a packet-switched WAN such as ARPAnet or a local area network (LAN) topology such as Ethernet.

TCP/IP to the Rescue

In 1973, work began on the TCP/IP protocol suite, a software-based set of networking protocols that allowed any system to connect to any other system, using any network topology. By 1978, IP version 4 (the same version that we use today) had been completed, although it would be another four years before the transition away from ARPAnet to IP would begin. Shortly thereafter, the University of California at Berkeley also began bundling TCP/IP with their freely distributed version of Unix, which was a widely used operating system in the research community.

The introduction and wide-scale deployment of TCP/IP represented a major ground-shift in computer networking. Until the introduction of TCP/IP, almost every other network topology required that hardware-based network nodes send traffic to a central host for processing, with the central host delivering the data to the destination node on behalf of the sender. For example, Figure 1-1 shows a host-centric networking architecture. In this model, devices are attached to a centralized system that coordinates all network traffic. A user at a terminal could not even send a screen of text to a printer without first sending the data to the central host, which would parse the data and eventually send it to the printer for printing.

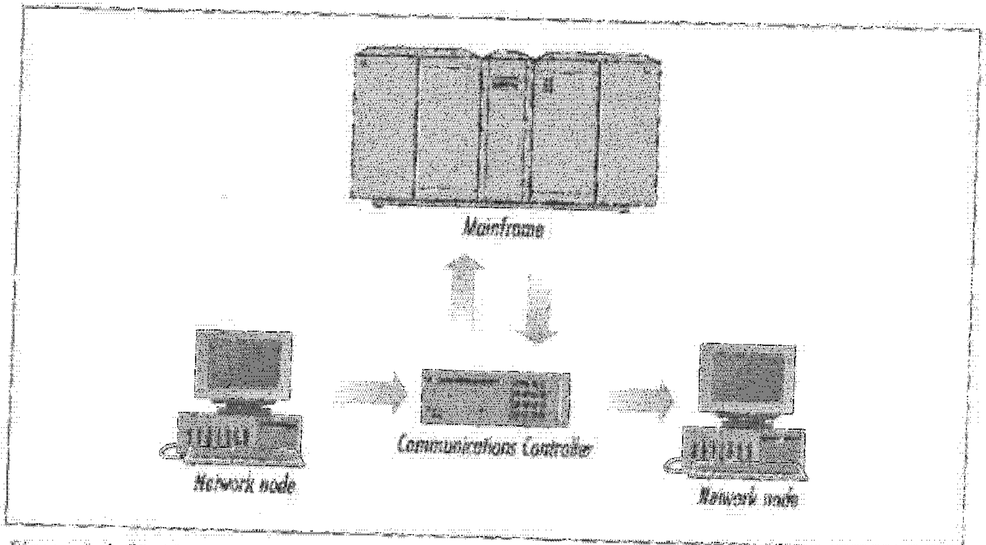

Figure 1-1. Host-centric networking

But with TCP/IP, each network device was treated as a fully functional, self-aware network end-point, capable of communicating with any other device directly, without having to talk to a central host first. IP networks are almost anarchic, with every device acting as a self-aware, autonomous unit, responsible for its own network services, as illustrated in Figure 1-2.

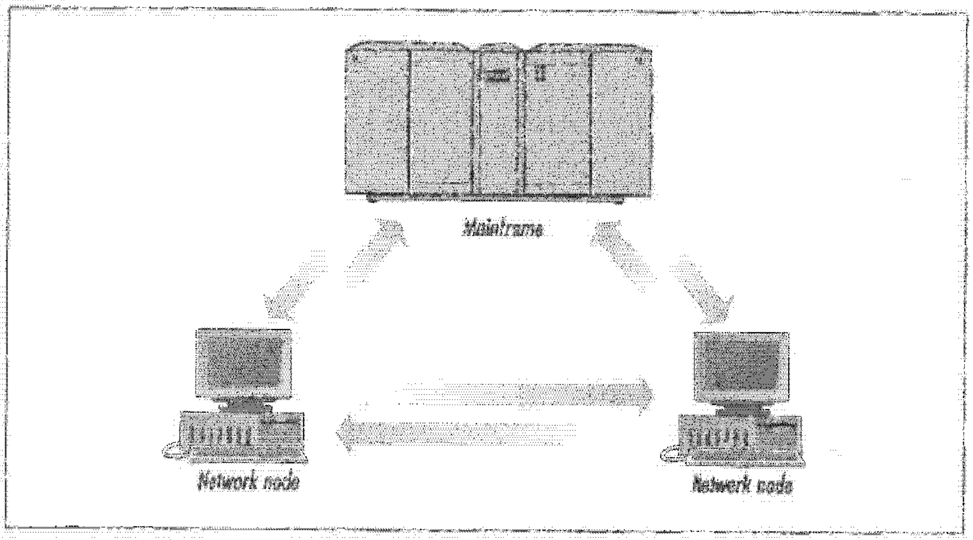

Figure 1-2. Node-centric networking

This design allowed for application- and resource-sharing on a national scale, since a top-down model simply would not work with millions of widely distributed devices. In addition, this design also provided reliability in case any part of the network was damaged, since a host-based model would simply stop functioning if the central host was destroyed or disabled.

The Internet Today

Over time, the ARPAnet evolved into an open "network-of-networks" using TCP/IP, with educational, commercial, and other organizations connected to each other through an interwoven mesh of networks. Today this type of mesh architecture is far less common, replaced by a much more structured hierarchy.

Rather than organizations connecting to each other directly, most organizations now connect to a local network access provider who routes network traffic upwards and outwards to other end-point networks.

Generally speaking, there are only a handful of top-level Internet Service Providers (ISPs), each of which provide major interconnection services around the country or globe. Most of these firms are telecommunications companies that specialize

in large-scale networking (such as long-distance providers like MCI WorldCom and Sprint).

Below these top-level carriers are local or regional access providers who offer regional access and lower-speed connection services to end users directly (these mid-level carriers are sometimes referred to as Internet Access Providers, or "IAPs"). This design is represented in Figure 1-3.

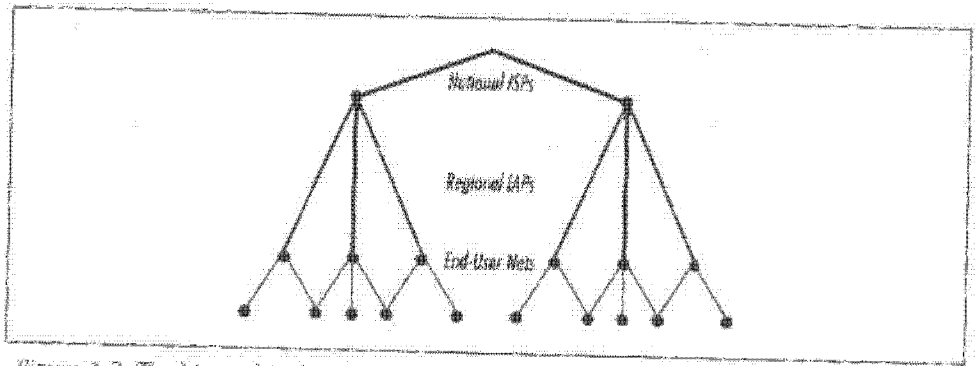

Figure 1-3. The hierarchical architecture of the Internet

Visually, the Internet can be thought of as a few major networking companies who provide large-scale "backbone" services around the world, followed by a large number of secondary providers that resell bandwidth on those networks. At the end of the line are the end-leaf organizations that actually generate the traffic that crosses these networks.

The Internet, Defined

Simply having a lot of interconnected networks does not by itself mean that you have the "Internet." To "internet" (with a lowercase "i") means to interconnect networks. You can create an internet of Macintosh networks using AppleTalk and some routers, for example. The term "Internet" (with a capital "I") refers to the specific global network of TCP/IP-based systems, originally consisting of ARPAnet and the other research networks.

There have been lots of private and public networks that have offered a multi-layer network design (private SNA* networks from the 1980s are a good example of this). Therefore, the Internet in particular is a collection of networks that support host-to-host communications using TCP/IP protocols.

* SNA stands for Systems Network Architecture, a proprietary IBM networking protocol.

Under this definition, the network is made up of intelligent end-point systems that are self-deterministic, allowing each end-point system to communicate with any host it chooses. Rather than being a network where communications are controlled by a central authority (as found in many private networks), the Internet is specifically meant to be a collection of autonomous hosts that can communicate with each other freely.

This is an important distinction, and one that is often overlooked. For example, many of the private networks have offered mail-delivery services for their customers, allowing a user on one network to send email to another user on another network, but only by going through a predefined mail gateway service. Conversely, the Internet allows users to exchange mail directly, without going through a central politburo first. In this regard, the Internet is a collection of self-deterministic, autonomous hosts.

Having hosts communicate with each other directly is not enough to make the Internet, however. Many networks have offered users the ability to communicate directly with other hosts on those networks, and those networks have not been considered as parts of the Internet per se. For example, there have been many private DECnet networks that have offered this capability, and Novell offers a similar service using IPX today.

The last key criteria is that the Internet is a collection of networks that allows host-to-host communications through voluntary adherence to open protocols and procedures defined by Internet standards. Therefore, in order for these networks to be parts of the Internet, they must also use Internet protocols and standards, allowing for vendor-neutral networking.

This is perhaps the most important part of the entire definition, since the use of consistent protocols and services is what allows the Internet to function at all. For example, it is not enough for a private network to allow users to send email messages to each other directly. Rather, those users must use the same protocols and services to exchange email messages, and those protocols must be defined as Internet standards.

TCP/IP's Architecture

A key part of understanding the distributed nature of TCP/IP is the realization that TCP/IP is a modular *family* of protocols, providing a wide range of highly segmented functions. TCP/IP is not a single monolithic protocol, but instead is a collection of protocols that range from application-specific functions like web browsing down to the low-level networking protocols like IP and TCP.